

CLAIMS:

Sub C1

1. A data processing system comprising a data storage device and a processor programmed to read data from, and write data to, said storage device, in which said storage device stores:

- multiple operation records each storing data relating to one or more historical operation involving at least one entity, each said operation record comprising data recording the operation, and data defining a date associated with the operation; and
- multiple entity records storing data indicating relationships between said entities, and each said relationship being associated with a historical period of validity.

2. The system of claim 1, wherein the processor is programmed to extract output data from a subset of said operation records, and to output said output data.

3. The system of claim 2, wherein the processor is programmed to select said subset by the steps of:

- inputting instructions defining one or more selected entities for which said output data relates; and
- selecting said subset based on both the dates stored in said operation records and the historical periods of validity associated with the selected entities.

4. The system of claim 3, wherein the processor is programmed to select said subset to represent by the steps of:

inputting an analysis date;

for the selected entities, selecting the entity relationships which have associated historical periods of validity within which said analysis date lies; and selecting said subset using those selected entity relationships.

5

5. The system of claim 4, wherein the processor is programmed to offer the current date as a date option, to permit analysis of operation records anterior to that date as if the current relationship between entities had previously existed.

10 6. The system of claim 4 or claim 5, wherein the processor is programmed to offer an anterior date as a date option, to permit analysis of operation records posterior to that date as if a historical relationship between entities still persisted.

15 7. The system of any of claims 3 to 6, wherein the processor is programmed to analyse each operation record in accordance with the relationships between entities which have associated historical periods of validity within which the date of that operation record lies.

20 8. The system of claim 1, wherein the processor is programmed to input a change from an existing said relationship between entities to a new said relationship.

9. The system of claim 8, wherein the processor is programmed, on such a change, to store an end date for the period of validity of the existing relationship; to create a record of the new relationship, and to store a start date therefor.

5 10. The system of claim 1, wherein the entity records comprise:
an entity record for each entity; and
an association record for each past or present relationship between a pair of said
entities;
each said entity record containing data representing its historical period of
10 validity.

11. The system of any preceding claim, wherein the entity records comprise a
hierarchical structure, in which at least a first entity record relates to a specific entity,
and a second to a more generic entity encompassing said specific entity, said entity
15 records including link data linking said first and second entity records whereby to allow
said processor to traverse said hierarchy.

12. The system of claim 11, wherein the entity records represent first and second
successive levels of hierarchy of an organisation.

20
13. The system of claim 11, wherein the entity records represent first and second
successive levels of hierarchy of a product family.

14. The system of claim 11 when dependant upon claim 3, wherein said processor is programmed to:

input a historical analysis period; and

determine, for said operation records within said period, if said operation records

5 relate to said selected entities throughout the whole of said period.

15. The system of claim 14, wherein, if said operation records do not span the whole of said period, for each selected said entity to which the operation records relate, the processor is programmed to determine from said entity records, a hierarchically higher

10 entity and to repeat said determination and, in the event that said operation records relate to said hierarchically higher entity throughout the whole of said period, to use said hierarchically higher entity instead of said selected entity in selecting said subset of operation records.

15 16. The system of any preceding claim in which said storage means contains multiple sets of said operation records, each said set comprising multiple said operation records, said sets relating to different classes of operations and said records within each set relating to different instances of the same type of operation.

20 17. The system of claim 16, in which each said operation record contains at least one variable data field storing a value of a measure from a range of possible said values for said measure.

18. The system of claim 16 or claim 17, in which said storage means further contains:

c) metadata comprising multiple operation definition records, each defining the format of records of a respective said set of operation records.

5

19. The system of claim 18 when dependant upon claim 17, in which each operation definition record indicates the units of said measure.

10 20. The system of claim 16 or claim 17, in which said storage means further contains:

c) metadata comprising multiple unit definition records, defining the relationship between different said units.

21. The system of claim 17, wherein the processor is programmed to:

15 input at least one measure derivable from said operation records, to be analysed; determine, for each said set of operation records, whether said measure can be derived therefrom; and,

where said measure could be derived from alternative said sets, select one of said sets.

20

22. The system of claim 21, wherein said selection is based on the relative sizes of said sets.

23. The system of claim 21 or claim 22, wherein said selection is based on the relative difficulty of deriving said measure from the data stored in the variable data fields of each of said sets.

5 24. The system of claim 17, wherein the processor is programmed to:
input at least one measure derivable from said operation records, to be analysed;
determine, for each said set of operation records, whether said measure can be
derived therefrom; and,
where necessary, derive said measure from a combination of a first value from a
10 variable data field of a record of a first set of operation records, and a second first value
from a variable data field of a record of a second set of operation records.

25. The system of claim 17, wherein the processor is programmed to:
input at least one measure derivable from said operation records, to be analysed;
15 determine, for each said set of operation records, whether said measure can be
derived therefrom; and,
where necessary, derive said measure from an aggregation of first values from
respective variable data fields of a plurality of records of a first set of operation records,
having dates spanning a predetermined input time interval.

20
26. The system of claim 1, wherein said operation records relate to respective
transactions between said entities.

27. The system of claim 26, wherein said transactions are sales, inventory, or purchase transactions.

28. The system of any preceding claim, wherein said processor is programmed to
5 load one or more new said operation records into said storage device.

29. The system of claim 28 when dependant upon claim 18, in which said processor
is programmed to determine whether said new operation records comply with said
metadata.

10 30. The system of claim 18, in which said processor is programmed to input said
metadata.

31. A data processing system, comprising:

15 processing means for generating a data model in accordance with a data
structure, the data model being adaptable to represent a change in the data structure; and
storage means for storing the data in accordance with the generated data model.

32. The data processing system of claim 31, wherein the stored data comprises
20 information representative of the time of change in the data structure.

33. The data processing system of claim 31 or 32, wherein the stored data
comprises:

transaction data representative of one or more measures which are determined relative to one or more references;

reference data representative of said one or more references; and

metadata descriptive of the transaction data and the reference data.

5

34. The data processing system of claim 33, wherein the metadata defines hierarchical associations between classes of the reference data.

35. The data processing system of claim 33 or 34, wherein the stored data comprises a number of elements of reference data, each element of reference data comprising information which defines an association with one or more other elements of reference data.

36. The data processing system of claim 35, wherein each element of reference data
15 further comprises information representative of a first period of validity of a defined
association.

37. The data processing system of claim 36, wherein the information representative
of the first period of validity comprises a start date of validity and an end date of
validity.

38. The data processing system of any of claims 33 to 37, wherein the one or more measures each are associated with one or more units.

39. The data processing system of claim 38, wherein the associations between the one or more measures and the one or more units are associated with a second period of validity.

5

40. The data processing system of claim 39, wherein the second period of validity comprises a start date of validity and an end date of validity.

41. The data processing system of any of claims 33 to 40, wherein the stored data 10 comprises a number of items of transaction data, each item of transaction data being associated with a date of transaction.

42. The data processing system of any of claims 33 to 40, wherein the metadata defines associations between classes of reference data and the one or more measures, the 15 associations between the classes of reference data and the one or more measures being representative of classes of transaction data.

43. The data processing system of any preceding claim, further comprising:
first interface means for receiving data of any structure from a data source for 20 storage in the data processing system.

44. The data processing system of any preceding claim, further comprising:

second interface means for outputting data from the storage means in a required format.

45. A data processing system, comprising:

5 processing means for generating a data model representative of data of a first structure, and for adapting the data model to represent also data of a second structure; and

storage means for storing data in accordance with the data model.

10 46. The data processing system of claim 45, wherein the stored data includes information representative of the time of adaptation of the data model.

47. A data storage device storing a data structure comprising:

15 a) multiple operation records each storing data relating to one or more historical operation involving at least one entity, each said operation record comprising data recording the operation, and data defining a date associated with the operation; and
b) multiple entity records storing data indicating relationships between said entities, and each said relationship being associated with a historical period of validity.

20 48. A data processing system comprising a data storage device and a processor programmed to read data from, and write data to, said storage device, in which said storage device stores multiple operation records each storing data relating to one or more historical operation involving at least one entity; and multiple entity records

SEARCHED
INDEXED
SERIALIZED
FILED
SUB C2

storing data indicating relationships between said entities, wherein the entity records comprise a hierarchical structure, in which at least a first entity record relates to a specific entity, and a second to a more generic entity encompassing said specific entity, said entity records including link data linking said first and second entity records 5 whereby to allow said processor to traverse said hierarchy, said processor being arranged to generate output data by inputting instructions defining one or more selected entity dimensions across which said output data is to be distributed.

49. The system of claim 48, wherein, if all required said operation records do not 10 relate to entities of the dimension to which the operation records relate, the processor is programmed to determine, from said entity records, a hierarchically higher level entity dimension and to repeat said determination and, in the event that all required said operation records relate to said hierarchically higher level, to use said hierarchically higher entity instead of said selected entity in selecting said subset of operation records.

15

50. The system of claim 48, wherein the processor is programmed to: 20
input at least one measure derivable from said operation records, to be analysed; and determine, for each said set of operation records, whether said measure can be derived therefrom; and, where said measure could be derived from alternative said sets, select one of said sets.

51. A method for managing data in a data structure, comprising the steps of:

storing a data model in said data structure, said data model identifying a plurality of relationships among classes of entities;

storing a plurality of datasets in said data structure, each of said datasets including an identification of at least one of said classes of entities and at least one
5 measure; and

relating each of said datasets to a respective subset of relationships in said data model, at least first and second ones of said datasets being related to respective first and second different ones of said subsets.

10 52. A method according to claim 51, wherein at least one of the relationships in said first subset is inconsistent with one of the relationships in said second subset.

53. A method according to claim 51, wherein the relationships identified in at least one of said subsets are hierarchical.

15 54. A method according to claim 51, wherein each of said entity class relationships further has associated therewith a temporal period of validity, and wherein said step of relating each of said datasets to a respective subset of relationships in said data model comprises the step of associating with each of said datasets a respective temporal period
20 of coverage.

55. A method according to claim 54, wherein the subset of relationships related to each given one of said datasets is defined to exclude each relationship in said data

model whose temporal period of validity does not overlap with the temporal period of coverage of the given dataset.

56. A method according to claim 51, further comprising the step of returning data
5 from said data structure in response to a query.

57. A method according to claim 51, further comprising the steps of:
determining, in response to a query requiring data values for a specified measure
distributed across entities in a specified class of entities, that not all datasets in said data
10 structure which are required to satisfy said query are broken down according to entities
in said specified class of entities; and

satisfying said query from datasets in said data structure at a level of granularity
of hierarchically higher class of entities which includes said specified class of entities.

15 58. A method according to claim 51, further comprising the steps of:
identifying, in response to a query, a superset of said datasets which together are
more than sufficient to satisfy said query;
selecting, in accordance with a predetermined optimisation function, a proper
subset of the datasets in said superset, the datasets in said subset being sufficient to
20 satisfy said query; and
satisfying said query from datasets in said subset.

59. A method for managing data in a data structure, comprising the steps of storing a data model in said data structure, said data model identifying a plurality of relationships among classes of entities each of said entity class relationships having associated therewith a respective temporal period of validity; and

5 storing a plurality of datasets in said data structure, each of said datasets including an identification of at least one of said classes of entities and at least one measure, and further having associated therewith a respective temporal period of coverage.

10 60. A method according to claim 59, wherein the subset of the relationships in said data model whose temporal periods of validity include a first temporal position are hierarchical.

15 61. A method according to claim 59, wherein the subset of the relationships in said data model whose temporal periods of validity include a first temporal position differs from the subset of the relationships in said data model whose temporal periods of validity include a second temporal position different from said first temporal position.

20 62. A method according to claim 61, wherein the subset of the relationships in said data model whose temporal periods of validity include said first temporal position is consistent with the subset of the relationships in said data model whose temporal periods of validity include said second temporal position.

63. A method according to claim 59, further comprising the step of returning data from said data structure in response to a query.

64. A method for extracting data from a data structure, said data structure including
5 a data model identifying a plurality of relationships among classes of entities, a first subset of said relationships having a hierarchical organisation, said data structure further including a plurality of datasets each including an identification of at least one entity and a data value for at least one measure, comprising the steps of:

10 receiving a query requiring data values for a specified measure distributed across entities in a specified class of entities;

determining that not all datasets in said data structure which are required to satisfy said query are broken down according to entities in said specified class of entities;

15 determining a hierarchically higher class of entities which includes said specified class of entities, said query being satisfiable from datasets in said data structure at the level of granularity of said hierarchically higher class of entities; and

satisfying said query from datasets in said data structure at the level of granularity of said hierarchically class of entities.

20 65. A method according to claim 64, wherein said step of receiving a query requiring data values for a specified measure comprises the step of receiving a query requesting information derivable from said data values for said specified measure.

66. A method according to claim 64, wherein said data model includes a plurality of subsets of said relationships among classes of entities including said first set of relationships, each of said subset having a respective hierarchical organisation, each of said datasets being related to a respective one of said subsets,

5 and wherein said step of determining that not all datasets in said data structure which are required to satisfy said query are broken down according to entities in said specified class of entities comprises the step of determining that a first one of the datasets in said data structure which are required to satisfy said query is related to said first subset of relationships whereas a second one of the datasets in said data structure 10 which are required to satisfy said query is related to a second one of said subsets of relationships different from said first subset of relationships, and first subset of relationships including a relationship involving said specified class of entities but said second subset of relationships not including any relationship involving said specified class of entities.

15 67. A method according to claim 66, wherein said first and second subsets of relationships share a sub-hierarchy which reaches down to the level of granularity of said hierarchically higher class of entities.

20 68. A method according to claim 66, wherein each of the relationships in said data model further has associated therewith a temporal period of validity, and each of said datasets has associated therewith a respective temporal period of coverage.

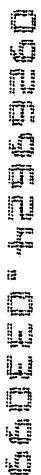
69. A method according to claim 68, wherein each of said datasets is related to a respective one of said relationship subsets at least in part by the temporal coverage of the dataset relative to the temporal period of validity of the entity class relationships in the respective subset.

5

70. A method according to claim 68, wherein said step of determining that a first one of the datasets in said data structure which are required to satisfy said query is related to said first subset of relationships whereas a second one of the datasets in said data structure which are required to satisfy said query is related to a second one of said

10 subsets of relationships different from said first subset of relationships, said first subset of relationships including a relationship involving said specified class of entities but said second subset of relationships not including any relationship involving said specified class of entities, includes the step of determining that a relationship exists in said data structure between said specified class of entities and a higher class of entities, 15 said relationship having a period of validity which overlaps with the temporal period of coverage of said first dataset but excludes the temporal period of coverage of said second dataset.

71. A method according to claim 64, wherein said step of satisfying said query from 20 datasets in said data structure at the level of granularity of said hierarchically higher class of entities, comprises the step of selecting datasets from said data structure which identify said hierarchically higher class of entities.

SEARCHED 

72. A method according to claim 64, wherein said step of satisfying said query from datasets in said data structure at the level of granularity of said hierarchically higher class of entities, comprises the steps of:

selecting the group of datasets from said data structure which identify classes of
5 entities which are hierarchically below said hierarchically higher class of entities; and
aggregating data values from the datasets in said group up to said hierarchically
higher class of entities.

73. A method according to claim 64, wherein said step of satisfying said query
10 comprises the steps of:

identifying, in response to said query, a superset of said datasets which together
are more than sufficient to satisfy said query at the granularity of said hierarchically
higher class of entities;

15 selecting, in accordance with a predetermined optimisation function, a proper
subset of the datasets in said superset, the datasets in said proper subset being both
necessary and sufficient to satisfy said query; and
satisfying said query from datasets in said proper subset.

74. A method for extracting data from a data structure, said data structure including
20 a plurality of datasets each including an identification of at least one entity in a
hierarchically defined class of entities, each of said datasets further identifying a data
value for at least one measure, comprising the steps of:

receiving a query requiring data values for a specified measure distributed across entities in a specified class of entities;

determining that not all datasets in said data structures which are required to satisfy said query are broken down according to entities in said specified class of 5 entities;

determining a hierarchically higher class of entities which includes said specified class of entities, said query being satisfiable from datasets in said data structure at the level of granularity of said hierarchically higher class of entities; and

10 satisfying said query from datasets in said data structures at the level of granularity of said hierarchically higher class of entities.

75. A method according to claim 74, wherein said step of determining that not all datasets in said data structure which are required to satisfy said query are broken down according to entities in said specified class of entities, comprises the step of determining 15 that none of the datasets in said data structures which are required to satisfy said query are broken down according to entities in said specified class of entities.

76. A method for extracting data from a data structure, said data structure including a plurality of datasets, comprising the steps of:

20 identifying, in response to a query, a superset of said datasets which together are more than sufficient to satisfy said query;

selecting, in accordance with a predetermined optimisation function, a proper subset of the datasets in said superset, the datasets in said subset being sufficient to satisfy said query; and

satisfying said query from datasets in said subset.

5

77. A method according to claim 76, wherein said query requires data values from a query set of at least one measure, and wherein said step of selecting a proper subset of the datasets in said superset, the datasets in said subset of being sufficient to satisfy said query, comprises the step of selecting a proper subset of the datasets in said superset, the datasets in which subset are sufficient to allow all of the required measures in said query set to be sourced.

10 78. A method according to claim 76, wherein said query cannot be satisfied from fewer than all of the datasets in said subset.

15

79. A method according to claim 76, wherein datasets in said data structure each include an identification of at least one entity in a hierarchically defined class of entities, and wherein said query requires data values for a specified measure distributed across entities in a specified class of entities, and wherein said step of identifying comprises the steps of:

20 determining a hierarchically higher class of entities which includes said specified class of entities, said query being satisfiable from datasets in said data structure at the level of granularity of said hierarchically higher class of entities; and

identifying a superset of said datasets which together are more than sufficient to satisfy said query at the granularity of said hierarchically higher class of entities.

80. A method according to claim 79, further comprising the step of determining that
5 not all datasets in said data structure which are required to satisfy said query are broken
down according to entities in said specified class of entities.

81. A method according to claim 76, wherein said query requests data broken down
with a particular level of granularity, and wherein said step of satisfying said query
10 comprises the step of satisfying said query with data broken down with a level of
granularity that is coarser than said particular level of granularity.

82. A method according to claim 76, wherein said step of selecting, in accordance
with a predetermined optimisation function, a proper subset of the datasets in said
15 superset, comprises the step of selecting, from all proper subsets of datasets in said
superset, the datasets in which subsets are both necessary and sufficient to satisfy said
query, a subset which has the lowest aggregate cost to satisfy said query.

add ex